

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 1 320 216 A1**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
18.06.2003 Bulletin 2003/25

(51) Int Cl.7: **H04L 12/18, H04L 12/26,
H04N 7/24**

(21) Application number: **01310356.9**

(22) Date of filing: **11.12.2001**

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR**
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: **The designation of the inventor has not yet been filed**

(74) Representative: **Lloyd, Barry George William et al
BT Group Legal Services,
Intellectual Property Department,
8th Floor, Holborn Centre,
120 Holborn
London EC1N 2TE (GB)**


(71) Applicant: **BRITISH TELECOMMUNICATIONS
public limited company
London EC1A 7AJ (GB)**

(54) **Method and device for multicast transmission**

(57) Apparatus for transmitting signals to video or other terminals 2 has coding means 5 for receiving the signals to generate a base signal stream B which is a coarse representation of the signal to be transmitted and at least one higher order stream E1, E2 containing finer detail of the signal not present in any lower order stream.

These streams are sent by transmitting means 6 to the terminals, but at least some of the terminals do not receive all of the higher order streams. Control means 7 receives signals x_k indicative of the data-receiving capacity of the terminals and calculates, based on the measured capacities, an adjustment for the coding means to make to the data rates of the streams.

Fig.1.



EP 1 320 216 A1

Description

[0001] The present invention is concerned with multicast distribution, that is, it envisages that a given signal feed (or recording) is to be distributed to a number of user terminals or receivers simultaneously. Commonly such material is distributed over digital telecommunications networks, for example, the internet or other internet protocol-based network.

[0002] The usual arrangement is that a source computer (the server) transmits the signals with a multicast address as the destination. A terminal (or client) that is to receive the signals corresponding to that multicast is said to join the corresponding multicast group: that is, the client's address is added to a list of group members, and routers within the network direct all packets which are addressed to that multicast address to all the member clients. In such circumstances the bit rate that can satisfactorily be received will vary from user to user, according, for example, to the type of connection he has, and moreover is likely to vary with time, depending on the prevailing level of network congestion.

[0003] One method of accommodating this is by the use of *layered* coding (also called *scalable* coding). Here the idea is that one generates two bit-streams, one - the base layer - carrying a low-quality signal, and a second - the enhancement layer - carrying at least part of the information that is missing from the base layer, so that a receiver that receives both layers can reconstruct a high-quality signal by combining the information received in both layers. More than two layers can be used, in which case each successive enhancement layer carries information that is denied to those receiving only the lower layer(s). A receiver can then be sent only as many layers as the currently available link between the transmitter and the receiver can comfortably handle. In the event of network congestion, the number of layers sent to the particular receiver can be reduced by ceasing transmission of the highest layer of those that were previously being sent. Layered coding is often used for the transmission of video signals, for audio signals, or both. In principle however it can be applied to any type of material for which a signal of reduced bit-rate (and, hence, in some sense, reduced resolution) is worthwhile receiving.

[0004] For discussion of layered video coding, see M. Ghanbari, *Video Coding - an introduction to standard codecs*, IEE, 1999, pp. 131-150.

[0005] For discussion of layered multicast, see S. McCanne, 'Receiver driven layered multicast', Proceedings of SIGCOMM 96, Stanford, CA, August 1996, and M. Nilsson, D. Dalby and J. O'Donnell, 'Layered Audio-visual Coding for Multicast Distribution on IP networks', Packet Video Workshop 2000, Cagliari, Italy.

[0006] A typical method of multicasting audiovisual material is that the audiovisual source multicasts layered audio and video over the IP network (internet). Each layer of audio and each layer of video is sent to a

separate multicast address. Each client makes its own decision of how many layers to receive, and joins the appropriate multicast groups. The content would have previously been advertised, using for example, SDP (session description protocol), which would include the bit rates and the addresses of the different layers.

[0007] One difficulty with layered multicasting is that it is non-optimum for a client terminal which has a data-receiving capacity which falls just short of that needed to support a particular layer, as it must then drop down to the layer below effectively "wasting" some of the capacity.

[0008] Another problem, identified (but not solved) by Nilsson et al in the above-referenced paper is that the client could not determine what network capacity was available between it and the server, other than by performing so-called join experiments, where the client randomly decides to join a higher layer, and then watches to see if the result is successful reception of the higher rate or packet loss across all layers. While this is reasonably satisfactory for a single client - some packet loss occurring when unsuccessful join experiments are conducted - it gets worse as more clients are involved, and make their own join experiments - as clients unaware of the join experiment may be affected by the packet loss resulting from an unsuccessful join experiment, and may incorrectly conclude that their network capacity from the server has reduced and so must reduce the number of layers received.

[0009] According to one aspect of the present invention there is provided method of transmitting signals to a plurality of terminals comprising:

a) coding the signals to generate a base signal stream which is a coarse representation of the signal to be transmitted and at least one higher order stream, where the or each higher order stream contains finer detail of the signal not present in any lower order stream;

b) transmitting the streams to the terminals, where at least some of the terminals do not receive all of the higher order streams;

c) receiving signals indicative of the data-receiving capacity of the terminals;

d) calculating, based on the measured capacities, an adjustment to be made to the data rates of the streams.

[0010] In another aspect, the invention provides an apparatus for transmitting signals to a plurality of terminals comprising:

a) coding means for receiving the signals to generate a base signal stream which is a coarse representation of the signal to be transmitted and at least

one higher order stream, where the or each higher order stream contains finer detail of the signal not present in any lower order stream;

b) means for transmitting the streams to the terminals, whereby at least some of the terminals do not receive all of the higher order streams;

c) control means for receiving signals indicative of the data-receiving capacity of the terminals and calculating, based on the measured capacities, an adjustment to be made to the data rates of the streams;

the coding means being responsive to control signals from the control means to adjust the rates of the streams.

[0011] In a further aspect, the invention provides a method of operating a telecommunications system in which a server transmits data to a terminal, wherein the method includes measuring the reception capacity of the terminal by:

(a) temporarily increasing the data rate of the data transmitted by the terminal, and

(b) measuring the time taken to receive the data transmitted at the increased rate.

[0012] Other aspects of the invention are set out in the claims.

[0013] Some embodiments of the invention will now be described, by way of example, with reference to the accompanying drawings, in which:

Figure 1 is a block diagram of a multicast video distribution system; and

Figure 2 is a block diagram of a layered video transcoder suitable for use in the system of Figure 1.

[0014] The system now to be described is for multicast distribution of layered video, but, as discussed above, can equally lend itself to the distribution of audio signals or other signals for which layered coding is practicable. It receives a video feed at an input 1, in compressed digital form - for example according to the MPEG-2 standard. The pictures which it represents are to be distributed to a number of user terminals 2. Each terminal is connected via some connection (perhaps via a telephone network or ADSL link) to one of a number of service provider points 3 which provide access to a telecommunications network (such as the internet) indicated at 4. Video signals arriving at the input 1 are coded by a layered video transcoder 5 which has a base layer output B, and first and second enhancement layer outputs E1, E2. These outputs are supplied to a network interface 6 connected to the network 4.

[0015] In a conventional arrangement of this kind, the layers each operate at a fixed bit-rate. Each client terminal determines how many layers it may receive and joins the corresponding multicast groups accordingly. For example if each layer operates at 64 kbit/s and the bit rate available for the particular terminal is 160 kbit/s, then the maximum number of layers is two, and the client terminal joins the lower two multicast groups only, so that the terminal is sent only the base layer B and the first enhancement layer E1. If the network conditions change, this decision may be modified so that the terminal switches to receiving only the base layer, or to receiving all three layers.

[0016] It will be observed that such a system does not make for the best picture quality available: in the case of the numerical example quoted above, the terminal, although capable of receiving at 160 kbit/s, in fact is receiving at only 128 kbit/s. There is also the drawback that any terminal whose currently available bit-rate falls below the rate needed for the base layer will suffer complete loss of signal until conditions improve once again.

[0017] In the system now proposed, the transcoder 5 has control inputs c_B , c_1 , c_2 by means of which the bit rate produced by the transcoder in each layer may be adjusted. A control unit 7 serves to receive, from the interface 6, signals x_k indicating the current bit-rate capacity available for each of the terminals currently connected, and periodically makes a decision as to any changes to the bit rates to be employed by each layer and sends control signals to the transcoder control inputs to adjust, if necessary, the bit rate in one or more layers.

The details of implementation of this process are

(i) to ascertain the rate x_k at which each terminal k ($k = 1..K$) is able to receive data over the network;

(ii) the algorithm to be used by the control unit 7 to determine the layer rates;

(iii) control of the transcoder to operate at the new rates; and

(iv) the process whereby the client terminals, following an adjustment to the layer rates, may change their decisions as to which layers to receive.

[0018] Considering first the available capacity, a first example of this may be implemented as follows.

[0019] The following example illustrates the principle of operation. Assume the server is sending a single stream at 10 packets per second. Normally it would try to space these evenly every 100ms. But if instead it sent 50 packets immediately, and then did nothing for the remainder of a 5 second period, before repeating the process for the next 50 packets etc, then assuming the client is capable of receiving at this average rate of 10 packets per second, then the time taken to receive this burst of

50 packets would give an indication of the throughput. So if they arrived in 2 seconds, it could conclude that it had the network capacity to receive an average of 25 (same sized) packets per second.

[0020] The impact on core network resources could be minimised by sending the higher layers in the gaps between bursts of packets in the lower layers, so that its overall output was a fairly constant packet rate. Note that for this technique to work, the server does not need to send packets back to back, but simply closer together in time than the client is capable of receiving.

[0021] Thus, the principle is

(a) the server transmits packets at an average rate of p packets per second.

(b) normally these packets are transmitted fairly regularly.

(c) periodically, (perhaps every 60 seconds) the server transmits in burst mode for a period of time T (e.g. 5 seconds), comprising T_B (e.g. 1) seconds during which it transmits packets at a rate of pT/T_B per second followed by $T - T_B$ seconds during which it transmits no packets. The burst packets are labelled with a marker to indicate that they belong to such a burst.

(d) the client terminal (which is programmed with, or informed of, the parameters p , T , T_B) upon recognition of a burst packet marker, starts a timer which is stopped when all pT packets with this marker have been received.

(e) if the timer count is T_R seconds, the client calculates the value $x_k = (pT - 1)/T_R$ (to allow for fact that $pT - 1$ packet intervals are timed, but pT packets are received).

(f) the client transmits this value x_k to the server.

[0022] Figure 1 illustrates this process in that the values x_k are shown as being received from the terminals via the network and the interface 6.

[0023] Alternatively, assuming a protocol involving the acknowledgment of received packet is in use, the measurement process could be performed at the server, which would measure the time duration between the first and last receipt by the server of acknowledgements of the packets of a particular burst.

[0024] It should be noted that this method of capacity measurement has uses beyond the present context of rate control of multistat streams, and may be used in other situations in which the data receiving capacity of a terminal needs to be measured.

[0025] In a simplified system, the available capacity could be input manually by a user to the terminals, based on the user's knowledge that the terminal is con-

nected via a particular type of link - e.g. a PSTN connection at 56 kbit/s or an ADSL connection at 1.5 Mbit/s. This is non-optimum, as it takes no account of variations in network congestion, but it would nevertheless represent a usable system.

[0026] Turning now to the control algorithm, the function of the control unit is

- to receive from each terminal a message stating the bit-rate currently available to it;
- to determine any necessary adjustments to the bit-rates of the layers.

[0027] In the following description, references to the bit-rate of a layer mean the bit-rate of that layer alone, whilst the total bit rate of a layer is used to mean the bit-rate needed by a terminal using it, that is to say, the sum of the bit rates of that layer and all lower layers. Terminals receiving the base layer only are referred to as Group 0 terminals, terminals currently receiving only the base layer and layer E1 are referred to as Group 1 terminals, whilst terminals receiving all three layers are referred to as Group 2. References in the following description to a terminal receiving a particular layer imply, of course, that that terminal is also receiving all of the lower layers too.

[0028] The algorithm to be employed by the control unit is not simple to devise. For example, if a Group 1 terminal reports a capability of 150 kbit/s and the current total bit rate of layer E1 is 128 kbit/s (say 64 kbit/s for the base layer and 64 kbit/s for layer E1), then manifestly the correct response may well not be simply to increase the total bit-rate of layer E1 to 150 kbit/s, since this may not be within the capabilities of other Group 1 terminals, which would then be forced to drop back to the base layer. Also there is an interdependence of the layers in the sense that a more appropriate response to a Group 1 terminal reporting surplus capacity may under some circumstances be to lower the total bit-rate of layer E2 so that that terminal may then switch to Group 2. Also it makes no sense to adjust the bit rate such that one layer ends up having a zero or trivially small bit rate.

[0029] Two possible algorithms will now be described. In each case, the reported bit rate capability of terminal k of a total of K terminals is denoted by x_k .

a) Given the distribution of bit rate capabilities x_k ,

- the lowest layer bit rate a_0 is chosen to be the lower of a fixed minimum value a_{0min} and a value x_{p0} that exceeds the x_k of no more than a fixed percentage P_0 (perhaps 0.1%) of the terminals, but may be increased to the lowest value of x_k if this is greater, i.e.

$$a_0 = \text{Max}\{\text{Min}\{x_k\}, \text{Min}\{a_{0min}, x_{p0}\}\}$$

- the other bit rates are chosen such that, of those terminals having x_k greater than or equal to x_0 , two thirds of them have values of x_k greater than or equal to $(a_0 + a_1)$ and one third have values of x_k greater than or equal to $(a_0 + a_1 + a_2)$.

b) Another approach is derived from the idea that users are charged for the number of bytes or packets received, and that the server wants to maximise revenue.

[0030] In simple terms, the rates a_0, a_1 and $a_2 \dots$ are chosen so that the sum, over all connected clients, of individual unused bit rate is minimised. For a given terminal k , with capacity x_k , between a_0 and $a_0 + a_1$, (i.e. $a_0 < x_k < a_0 + a_1$) the unused bit rate would be $x_k - a_0$.

$$\text{Let } n_s = \sum_{i=0}^{s-1} a_i$$

(i.e. n_s is the total rate of s layers)

[0031] Suppose that terminal k operates at a bit rate r_k which is the largest $n_{s \leq x_k}$

[0032] Choose a_i ($i = 0 \dots l$) such as to minimise the unused capacity

$$U = \sum_{k=1}^K (x_k - r_k) = \sum_{k=1}^K x_k - \sum_{k=1}^K r_k$$

[0033] Since, however, the x_k are given, the task is to maximise the term

$$V = \sum_{k=1}^K r_k$$

[0034] Given that r_k is not a continuous function of a_i , an elegant solution of this maximisation problem seems improbable and thus we propose that the control unit is programmed to solve the problem by evaluating V for all possible values of a_i within defined ranges, subject to some quantisation on the choices for a_i - perhaps multiples of 1 kbit/s.

[0035] Having determined the bit rate for each layer, the control unit communicates the new rates to the transcoder

[0036] Turning now to the transcoder, a block diagram of a layered transcoder with adjustable rates is shown in Figure 2. This is a layered, SNR-scalable transcoder operating in accordance with the MPEG-2 standard. An incoming MPEG-2 signal at an input 10 firstly passes through a decoder 11 for variable-length codes, followed

by an inverse quantiser 12 matched to the quantiser of the original encoder. The coefficients output from the inverse quantiser 12, after addition at 101 of a correction signal (to be described below) are requantised in a coarse quantiser 102 and variable-length coded at 103 to provide a base-layer signal which is then input to a buffer 104.

[0037] A subtractor 105 forms the difference between the coefficients at the output of the inverse quantiser 12, and those obtained via an inverse quantiser 106 from the base layer signal at the output of the quantiser 102. This difference signal, which represents the information now lacking from the first layer, forms the input to the next layer.

[0038] Drift correction (if required) is provided as follows. The difference signal from the subtractor 105 is converted from the frequency domain to the pixel domain by an inverse discrete cosine transform (DCT) unit 107 and then accumulated in a loop comprising an adder 108, a frame store 109 and a motion compensation unit 110 which receives motion vectors extracted (by means not shown) from the incoming signal. Note that the frame store 109 is shown as multiple stores, which are needed in systems such as MPEG which use an irregular inter-frame prediction sequence. The arrangements for switching between the stores are as in a conventional MPEG coder.

[0039] This accumulated drift is then converted to the frequency domain by a DCT unit 111 and added as a correction to the next frame via the adder 101.

[0040] A control unit 112 serves to sense the state of fullness of the buffer and to influence the rate of generation of data in such a manner as to ensure that the buffer neither overflows or becomes empty; this it does by controlling the step-size of the quantiser 102, or by causing frame-dropping by a switch 113, or both. Commonly in systems of this type, data would be read out of the buffer at a constant bit-rate. Here, however the buffer receives at a control input 114 the aforementioned control signal C_B which serves to set the rate of a readout clock to the rate a_0 indicated by C_B .

[0041] In principle, the control unit 112 could use a conventional buffer control algorithm - that is, the rate control signal C_B simply changes the readout rate and the control unit tracks this by monitoring the state of buffer fullness. In practice, however, it may be preferred to use the control signal C_B to influence the transcoder more directly, for example simple rate control strategy is to set a picture dropping threshold proportional to the bit rate (such as 4 picture periods x bit rate), and to select the quantiser to be proportional to the buffer fill, so that the finest quantiser is used when the buffer is empty and the coarsest one is used when the level is the picture dropping threshold.

[0042] The first enhancement layer is constructed in identical manner to the base layer, and receives the output from the subtractor 105 at adder 201 - components 201 to 214 are identical to components 101 to 114 of the

base layer. The control input 214 receives the first enhancement layer control signal C_1 . Similarly, the second enhancement layer has components 301 to 313 and control input 314 receiving the control signal C_2 . Naturally, further layers could be added, on the same pattern, if desired.

[0043] This and other layered transcoders are described in our co-pending U.K. patent application no. ***** (Applicants' reference A30104) filed on the same date as the present application.

[0044] In order to create the bursty transmission to allow the client to measure the network capacity, the readout rate on each layer need periodically to be temporarily increased. For this purpose, control signals B_B , B_I and B_2 each consisting of a differently-phased 1-second pulse every 60 seconds are generated by the control unit 7 and supplied to the respective buffers 104, 204, 304 to increase the readout rate. Note that if this approach is adopted, to avoid the buffer control being affected it is necessary for the control units 112, 212, 312 to calculate the fullness that the relevant buffer would have had with uniform emptying. Alternatively, a second buffer could be added to each layer, with uniform packet transfer (at rate a_i) from the first, and bursty transmission from the second into the network.

[0045] Finally, there remains the question of how the clients make if necessary, a revised decision as to which layers they subscribe to. One possibility would be for them to follow the existing practice of "join experiments" as described in the introduction. A more attractive option, in the present context, is to make use of the fact that the data rate x_k available to the client has already been determined. Thus the server notifies the clients of the new bit rates of each layer (via an SDP advertisement as mentioned earlier). The client then determines how many layers it can receive within the limit corresponding to x_k for that terminal, and if necessary joins or leaves multicast groups so that it is sent those layers and no others.

[0046] Note: if the available data rates x_k were to be measured by the server, then the server would need to communicate this information to the respective terminals. Alternatively, the server could perform the same calculation and instruct the client which groups to subscribe to.

[0047] Other forms of layered transcoder can also be used, including those with spatial scaling. Indeed, the control mechanisms envisaged here could (with uncompressed video input) be used in a system using layered encoders (as opposed to transcoders), or indeed to encoders and transcoders that produce hierarchies of independent streams, rather than adopting a layered approach.

Claims

1. A method of transmitting signals to a plurality of ter-

minals comprising:

- a) coding the signals to generate a base signal stream which is a coarse representation of the signal to be transmitted and at least one higher order stream, where the or each higher order stream contains finer detail of the signal not present in any lower order stream;
- b) transmitting the streams to the terminals, where at least some of the terminals do not receive all of the higher order streams;
- c) receiving signals indicative of the data-receiving capacity of the terminals;
- d) calculating, based on the measured capacities, an adjustment to be made to the data rates of the streams.

2. A method according to claim 1 in which the or each higher order stream lacks information contained in a lower order stream, in which each terminal receiving a higher order stream receives also all lower order streams.
3. A method according to claim 1 or 2 in which the streams are generated by a transcoder.
4. A method according to claim 1, 2 or 3 in which the signals indicative of the data-receiving capacity are manually input.
5. A method according to claim 1, 2 or 3 in which the signals indicative of the data-receiving capacity are obtained by measuring the said capacity.
6. A method according to claim 5 in which said measurement is performed by temporarily increasing the data rate of a transmitted stream and measuring the rate at which it is received by the terminal.
7. A method according to claim 5 in which said measurement is performed by temporarily increasing the data rate of a transmitted stream and measuring the rate at which acknowledgements thereof are received from the terminal.
8. A method according to any one of the preceding claims in which the calculation of data rates to be used for transmission is performed by selecting rates such that the numbers of terminals capable of receiving each level of quality shall be in predetermined proportions.
9. A method according to any one of claims 1 to 7 in which the calculation of data rates to be used for transmission is performed by calculating the difference between the total data-receiving capacity of the terminals and the total of the capacity required for each terminal to receive the highest quality sig-

nal requiring a bit rate less than or equal to its capacity, and calculating an adjustment to be made to the data rates of the streams such as to tend to reduce that difference.

10. An apparatus for transmitting signals to a plurality of terminals comprising:

a) coding means for receiving the signals to generate a base signal stream which is a coarse representation of the signal to be transmitted and at least one higher order stream, where the or each higher order stream contains finer detail of the signal not present in any lower order stream;
 b) means for transmitting the streams to the terminals, whereby at least some of the terminals do not receive all of the higher order streams;
 c) control means for receiving signals indicative of the data-receiving capacity of the terminals and calculating, based on the measured capacities, an adjustment to be made to the data rates of the streams;
 the coding means being responsive to control signals from the control means to adjust the rates of the streams.

11. An apparatus according to claim 10 in which the or each higher order stream lacks information contained in a lower order stream, in which each terminal receiving a higher order stream receives also all lower order streams.

12. An apparatus according to claim 10 or 11 in which the coding means is a transcoder.

13. An apparatus according to claim 10, 11 or 12 in which control means is operable to receive the signals indicative of the data-receiving capacity from the terminals.

14. An apparatus according to claim 13 including means operable to temporarily increase the data rate of a transmitted stream whereby the rate of its receipt may be measured by a terminal.

15. An apparatus according to claim 10, 11 or 12 including means operable to temporarily increase the data rate of a transmitted stream, and means measuring the rate at which acknowledgements thereof are received from the terminal.

16. An apparatus according to any one of claims 10 to 15 in which the control means is arranged in operation to calculate data rates to be used for transmission by selecting rates such that the numbers of terminals capable of receiving each level of quality shall be in predetermined proportions.

17. An apparatus according to any one of claims 10 to 15 in which the control means is arranged in operation to calculate data rates to be used for transmission by calculating the difference between the total data-receiving capacity of the terminals and the total of the capacity required for each terminal to receive the highest quality signal requiring a bit rate less than or equal to its capacity, and calculating an adjustment to be made to the data rates of the streams such as to tend to reduce that difference.

18. A method of operating a telecommunications system in which a server transmits data to a terminal, the method including measuring the reception capacity of the terminal by

a) temporarily increasing the data rate of the data transmitted by the server to the terminal, and
 b) measuring the time taken to receive the data transmitted at the increased rate.

19. A method according to Claim 18 in which the step of measuring the time taken is performed by the terminal.

20. A method according to Claim 18 in which the data are transmitted in packets, the terminal sends acknowledgements of receipt of the packets to the server, and the step of measuring the time taken is performed by the server measuring the time taken to receive the acknowledgements of the packets transmitted at the increased rate.

21. A transmission system comprising a server for transmitting data and

a terminal for receiving data, wherein the terminal includes means operable to temporarily increase the data rate of a transmitted stream, and the system includes means operable to measure the time taken to receive the data transmitted at the increased rate.





Fig.2.





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 31 0356

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	ZHAO SONGSHENG ET AL: "Dynamic quality of session control of real-time video multicast" INTELLIGENT PROCESSING SYSTEMS, 1997. ICIPS '97. 1997 IEEE INTERNATIONAL CONFERENCE ON BEIJING, CHINA 28-31 OCT. 1997, NEW YORK, NY, USA, IEEE, US, 28 October 1997 (1997-10-28), pages 1737-1741, XP010276324 ISBN: 0-7803-4253-4 * page 1738, left-hand column, line 46 - page 1740, left-hand column, last line *	1,2,4,5, 8,10,11, 13,16	H04L12/18 H04L12/26 H04N7/24
Y	-----	3,6,7, 12,14,15	
Y	RAMANUJAN R S ET AL: "Adaptive streaming of MPEG video over IP networks" LOCAL COMPUTER NETWORKS, 1997. PROCEEDINGS., 22ND ANNUAL CONFERENCE ON MINNEAPOLIS, MN, USA 2-5 NOV. 1997, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 2 November 1997 (1997-11-02), pages 398-409, XP010252445 ISBN: 0-8186-8141-1 * page 398, right-hand column, line 6 - page 399, left-hand column, line 8 * * page 401, left-hand column, line 34 - right-hand column, line 12 *	3,12	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
			H04L H04N
X	US 6 215 774 B1 (KNAUERHASE ROBERT C ET AL) 10 April 2001 (2001-04-10) * abstract *	18-21	
Y	-----	6,7,14, 15	
	----- -/--		
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 7 June 2002	Examiner Ströbeck, A.
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

EPO FORM 1503 03 82 (P/C/C01)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 31 0356

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	<p>VICKERS B J ET AL: "Adaptive multicast of multi-layered video: rate-based and credit-based approaches"</p> <p>INFOCOM '98. SEVENTEENTH ANNUAL JOINT CONFERENCE OF THE IEEE COMPUTER AND COMMUNICATIONS SOCIETIES. PROCEEDINGS. IEEE SAN FRANCISCO, CA, USA 29 MARCH-2 APRIL 1998, NEW YORK, NY, USA, IEEE, US, 29 March 1998 (1998-03-29), pages 1073-1083, XP010270328</p> <p>ISBN: 0-7803-4383-2</p> <p>* page 1073, right-hand column, line 8 - line 39 *</p> <p>* page 1074, right-hand column, line 43 - page 1076, right-hand column, last line *</p> <p>-----</p>	1,2,4,5,10,11,13,14	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 7 June 2002	Examiner Ströbeck, A.
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 503 03 02 (P01001)

